Verifying Web Applications: From Business Level Specifications to Automated Model-Based Testing

Christian Colombo Mark Micallef Mark Scerri

PEST Research Lab
Department of Computer Science
University of Malta, Malta

{christian.colombo | mark.micallef | msce0013}@um.edu.mt

One of reasons preventing a wider uptake of model-based testing in the industry is the difficulty which is encountered by developers when trying to think in terms of properties rather than linear specifications. A disparity has traditionally been perceived between the language spoken by customers who specify the system and the language required to construct models of that system. The dynamic nature of the specifications for commercial systems further aggravates this problem in that models would need to be rechecked after every specification change. In this paper, we propose an approach for converting specifications written in the commonly-used quasi-natural language Gherkin into models for use with a model-based testing tool. We have instantiated this approach using QuickCheck and demonstrate its applicability via a case study on the *eHealth* system, the national health portal for Maltese residents.

1 Introduction

In typical commercial contexts, modern software engineering is characterised by a strong customer focus and, as a result, fluctuating requirements. The industry has developed ways to adapt to this, most notably through Agile development processes [5] which have gained substantial popularity over the past decade. These processes start off by gathering requirements from the customer in the form of case scenarios and subsequently development is built around these scenarios culminating in their implementation and demonstration to the customer. This process, better known as behaviour-driven development [3], reorientates the development process so as to ground all work around business level customer specifications. Once the customer scenarios are implemented, the process tends to take on an iterative form whereby the customer is again asked to review the specification and possibly introduce further scenarios. To keep the specification language as simple as possible for non-technical customers, scenarios are typically specified in terms of a quasi-natural language called Gherkin [8]. In order to automate test execution, clauses from the Gherkin specification are associated with executable code using a technology such as Cucumber [6]. This executable code effectively turns the specification into an automated acceptance test which is used to test that the system adheres to the customer's specifications.

Whilst the automation achieved through the running of Gherkin specifications has significantly shortened the test execution process and made it repeatable, a significant limitation is that the scenarios are never executed in combination with one another: one would typically expect some classes of problems to arise through the interactions between a sequence of scenarios rather than within the context of individual scenarios. In this context, we present an approach to combine related scenarios into models which can be consumed by model-based testing tools. In turn, such tools can be used to explore the model, generating plausible sequences of scenarios that provide more thorough testing of the system for minimal extra effort. We show how this technique has been instantiated on the QuickCheck testing tool and applied to Malta's national health portal web application. This paper makes three main contributions:

A. Petrenko, H. Schlingloff (Eds.): Ninth Workshop on
Model-Based Testing (MBT 2014)
EPTCS 141, 2014, pp. 14–28, doi:10.4204/EPTCS.141.2

(i) A technique for automatically generating system models from existing test case scenarios written in the Gherkin language. These models can be consumed by model-based testing tools in order to generate test cases that extend normal testing and facilitate inter-scenario interaction. (ii) A proof-of-concept tool which applies the technique presented here to generate models for the QuickCheck test case generation tool. (iii) A case study on Malta's national health portal which demonstrates that the technique can be applied to a real-world system.

In order to control scope at this stage of our research, we constrain our work to the context of web applications. Whilst this is a plausible constraint due to the increased proliferation of web applications, it does not preclude the technique from being modified and applied to other contexts. The rest of the paper is organised as follows: Section 2 provides the reader with the background required to understand the work presented in this paper whilst Sections 3 and 4 explain the proposed technique and provides an overview of the instantiation respectively. This is followed by Section 5 whereby we discuss a case study designed to demonstrate the applicability of the technique. Finally, Section 6 discusses related work from the literature and Section 7 features our conclusions and opportunities for future work.

2 Background

Gherkin is currently an industry standard language enabling non-technical customers to express requirements in terms of case scenarios. Ensuring that the developed system adheres to these specifications is crucial for customer acceptance of the resultant system. By associating executable code to statements within Gherkin, developers can use tools such as Cucumber to automatically execute test scenarios on the underlying system. However, to enable the automated manipulation of web applications through a browser, mimicking the user, one would need a tool such as Selenium which provides the necessary interface to access a web application programmatically. The rest of this section provides the necessary detail on these technologies so as to enable the reader to understand the work presented here.

2.1 Gherkin

Gherkin [8] is a business readable, domain-specific language which lets users describe what a software system should do without specifying how it should do it. The language forms an integral part of *Behaviour Driven Development* [3], a development process which aims to make specification a ubiquitous language throughout a system's life cycle. Gherkin specifications take on a so-called *Given-When-Then* format as illustrated in Example 1 and the underlying concept is that a business user can specify what is expected from a system using a series of concrete scenarios. *Given* some precondition, *When* the user does something specific, *Then* some postcondition should materialise.

Example 1

```
1. Feature: Doctor's landing page
2. Scenario: Navigate to lab results page
3. Given I am on the doctors landing page
4. When I click on laboratory results
5. Then I should go to the lab results page
6. Scenario: Navigate back to Doctor's landing page
7. Given I am on the lab results page
8. When I click on the myHealth logo
9. Then I should go to the doctors landing page
```

In this example, the client has specified two requirement scenarios as part of a feature on an eHealth system: one dealing with the process of a doctor viewing the laboratory results and another one of a doctor navigating back to the landing page.

Developers then decompose each scenario into a number of development tasks and the work proceeds with all activity being a direct result of a particular customer-specified scenario. Furthermore, although the scenarios are written in a subset of natural language, technologies exist (e.g. Cucumber [6] and Specflow [14]) which enable developers to implement automated user acceptance tests for every step in a scenario. That is to say that (i) the user specification becomes executable as an automated test by the end of a development iteration; and (ii) any future scenarios which include a predefined step will automatically reuse test automation code, thus in effect building up a domain specific language for automated testing.

2.2 Cucumber

Cucumber [6] is a technology that enables developers to provide glue code which turns natural language specifications into executable automated tests. Although originally written in Ruby, the tool has been ported to a number of languages. It reads plain language text files called *feature files* (as in Example 1), extracts scenarios and executes each scenario against the system under test. This is achieved by means of so-called *step definitions*, a series of methods which are executed when a step in a feature file matches a particular pattern. For example for the statement `When I click on the myHealth logo`, the developer associates code which actually simulates a user clicking on the logo. Once all Gherkin statements have associated code which can be executed by Cucumber, the latter will iterate through each step, find a matching step definition and execute the method associated with it. Step definitions can implement any feature supported by the underlying language. In our case, we are interested in manipulating web browsers for automated web application testing; something which can be achieved by Selenium.

2.3 Selenium

Selenium [13] is a web browser automation framework which has become a de-facto industry standard when it comes to web test automation. It has been implemented in a number of popular programming languages and provides API-level access to various browsers. In essence, developers can write programs which launch a browser, simulate user behaviour within the browser and subsequently query the state of website elements to verify expected behaviour as illustrated in Example 2.

Example 2

```
public void testGoogle() {
    // Launch browser and visit www.google.com
     WebDriver browser = new FireFoxDriver();
     browser.open("http://www.google.com");
    // Search for ‘‘Hello World”
     browser.findElement(By.id("q")).sendKeys("Hello World");
     browser.findElement(By.id("btnG")).click();
    // Check that the search engine behaved appropriately
     assertEqual("Hello World - Google Search", browser.getTitle());
}
```

The above code snippet features source code written in Java which carries out a test on the Google search engine using the Firefox browser. The `WebDriver` *object in this example is provided by the Selenium framework.*

Combining Cucumber with Selenium can result in truly powerful automated test suites which provide repeatable and consistent regression testing capabilities. We argue that this combination can be leveraged even further by bringing model-based testing into the mix at minimal cost by constructing models from existing Gherkin specifications.

3 Proposed Technique

Starting with a number of Gherkin scenarios, the proposed technique combines characteristics inherent in the *Given-When-Then* notation with those of web applications in order to make educated inferences and consequently construct a model of a specified system. In order to construct a model representation of web applications, the specification language should allow us to:

1. Define the states of the system
2. Specify the starting state
3. Specify the transitions available at each state, including any preconditions, the action which causes the system to change state, and any postconditions

We argue that the *Given-When-Then* structure of Gherkin lends itself to providing this information since a single scenario in *Given-When-Then* form can be interpreted to represent a single transition between two states. That is to say that the *Given* step indicates an origin state (and possibly preconditions), the *When* state specifies actions for a transition to occur, and finally the *Then* step indicates the target state (possibly with postconditions). However, in order to facilitate the automated interpretation of scenarios in this manner, we propose that users of the language adopt three conventions taking into account the navigation-driven nature of web applications. To illustrate the conventions, we will be using Example 3 as a running example.

Example 3

```
 1. Feature: Patient management features
 2. Scenario: Navigate to lab results page
 3. Given I start on the "doctors landing page"
 4. And I have pending lab results
 5. When I click on laboratory results
 6. Then I should go to the "lab results page"
 7. And I should see my pending lab results
 8. Scenario: Navigate back to Doctor's landing page
 9. Given I am on the "lab results page"
10. When I click on the myHealth logo
11. Then I should go to the "doctors landing page"
```

The above Gherkin code consists of two scenarios which are authored according to the conventions proposed here. More details on the pattern of these scenarios are given below when describing the conventions.

Corresponding to the three points outlined about, we propose the conventions below:

Convention 1: Make states explicit. In order to identify states within a test scenario, we constrain the language such that *Given* steps take the form of ***Given I am on the "<web page >"***. Similarly, *Then* steps take the form of ***Then I should be on the "<web page >"***. Recall that we consider web pages within an application to represent states within the application's model. This convention, whereby *Given* and *Then* steps take a specific form which includes state names being surrounded by quotes, enables us to clearly identify states within the model. In Example 3, the states "doctors landing page" and "lab results page" can be extracted from lines 3, 6, 9 and 11. Upon closer inspection, one realises that line 3 takes on a slightly different form. This is outlined in the next convention.

Convention 2: Identify start states. In order to identify start states in a model, we propose that the convention on *Given* steps consist of two variations. The first is the ***Given I am on the "<web***

page >" outlined in *Convention 1* whereas the second would replace "I am" with "I start" resulting in ***Given I start on the "<web page >"***. This minor variation will identify the web page in question as a starting state in the system's model.

Convention 3: Identify actions, preconditions and postconditions. Finally, we identify the *When* construct as the action which takes the user from one page to another while the *And* steps which form part of *Given* or *Then* parts of scenarios will be leveraged to identify preconditions and postconditions respectively. That is to say that this additional information can be used to infer conditions which should be satisfied before a transition can fire, as well as conditions which should be satisfied after a transition has fired. Consider the first scenario in Example 3. The action taking the user from the "doctors landing page" to the "lab results page" is identified by the *When* clause: "clicking on laboratory results". Furthermore, in this case, the *Given* portion of the scenario has two lines. The first identifies the current state as the "doctors landing page" whilst the second sets up a precondition on the subsequent transaction. In order for "clicking on laboratory results" to be a valid action, first the condition that the doctor has pending lab results must be satisfied. Similarly the *Then* portion of the scenario identifies the "lab results page" as the target state of the scenario whilst line 7 states that the doctor should also see pending lab results within that page once the transition is complete.

If these conventions are adhered to, test scenarios can be interpreted as follows:

Given: The *Given* step is usually used to specify preconditions in a test scenario but in the context of model-based testing, it can be used to infer the current state of the system. In accordance to convention 2, a special adaptation of this step can also be used to specify the starting state of a model. Furthermore, any *And* steps immediately following the *Given* step will be treated as preconditions.

When: In normal use, the *When* step specifies actions for exercising the system under test in a scenario and retains this role in our technique such that the action specifies a transition from the current state to a target state.

Then: The *Then* step is usually utilised to specify a postcondition in a test scenario but in our technique we adapt this to specify the target state. Furthermore, any *And* steps immediately following the *Then* step will be treated as postconditions.

By processing multiple test scenarios and joining them at points where states from different scenarios have the same name, a model can now be constructed. This is demonstrated in Figure 1 whereby four generic test scenarios are combined to construct a single model. More formally, given a list of test scenarios written in Gherkin such that they adhere to the conventions discussed above, a model can be constructed using Algorithm 1. In short the algorithm processes the scenarios, identifies unique states and inserts the transitions including actions, preconditions and postconditions. Figure 2 shows a model constructed by applying this algorithm to Example 3.

In the next section, the approach described above is instantiated in the context of a particular model-based testing tool.

4 Instantiation

As a prerequisite to evaluating the approach, we developed an instantiation: namely stringing together three parts: (i) the algorithm translating the Gherkin scenarios into a model for a model-based testing

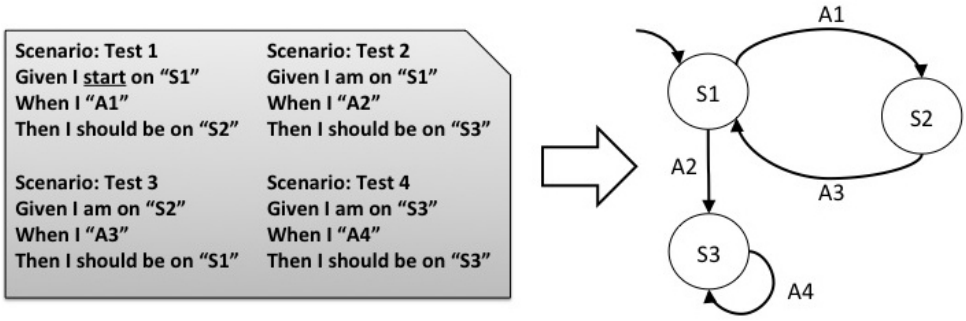

Figure 1: A depiction of how models can be constructed from multiple test scenarios.

Data: scenarios : List of Test Scenarios
Result: model : A finite state machine modelling the system under test.
model = new Model();
while *scenarios.hasMoreScenarios()* **do**
    scenario = scenarios.nextScenario();
    os = scenario.extractOriginState() ; `// According to Convention 1`
    ts = scenario.extractTargetState() ; `// According to Convention 1`
    if *(os != null* $\wedge$ *ts != null)* **then**
        model.addStateIfNew(os) ;
        model.checkIfStarting(os) ; `// May be marked as start state if Convention 2`
        model.addStateIfNew(ts) ;
        preconditions = scenario.extractPreconditions() ; `// According to Convention 3`
        postconditions = scenario.extractPostconditions() ; `// According to Convention 3`
        actions = scenario.extractActions() ; `// All actions from the When steps`
        if *(actions != null)* **then**
            transition = new Transition(preconditions, actions, postconditions);
            model.addTransition(os,ts,transition);
        else
            error(Unable to extract transition from test scenario);
        end
    else
        error(Unable to extract states from test scenario);
    end
    return model;
end

Algorithm 1: High level algorithm for generating models from business specifications.

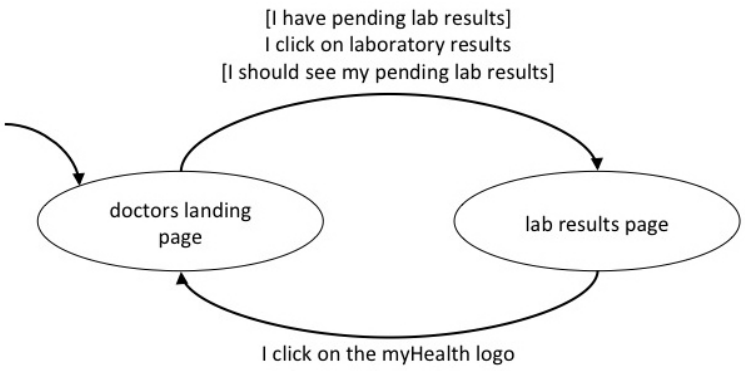
[I have pending lab results]
I click on laboratory results
[I should see my pending lab results]
doctors landing page
lab results page
I click on the myHealth logo

Figure 2: A model constructed by applying Algorithm 1 to Example 3.

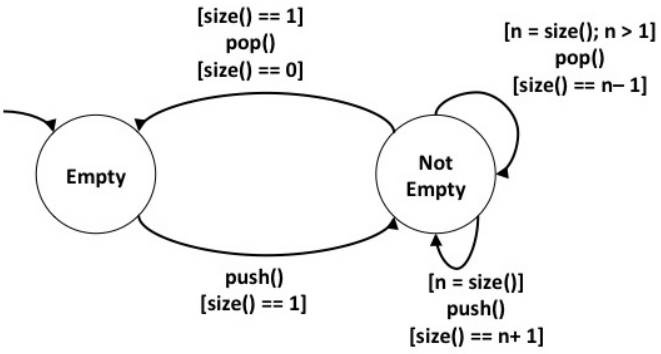

Figure 3: An example of a model specified in QuickCheck.

tool, (*ii*) the model-based testing tool which generates the test cases and detects test failure or success, and (*iii*) Selenium which enables the testing tool to interact directly with the web application.

While other model-based testing tools such as ModelJUnit[1] are perfectly valid alternatives, as a model-based testing tool, we selected QuickCheck [4] for the fact that it provides test case *shrinking* functionality whereby upon detecting a failure, it searches for smaller test cases which also generate the fault to aid debugging. Although initially developed as a tool for the Haskell programming language, it has been implemented on a number of platforms and for this work, we utilise the Erlang implementation [2]. QuickCheck is primarily a random test data generation tool but it also supports model-based testing. Given a model and a set of pre and postconditions, the tool is able to automatically generate test cases which are used to verify that the conditions hold in the context of different sequences of actions. Figure 3 illustrates an example QuickCheck model which represents valid sequences of operations involving a stack (top elements in square brackets represent preconditions while the bottom elements represent postconditions). Elements can be pushed or popped from the stack: Popping an element from the stack is not allowed when the stack is empty, otherwise the size should be decremented by one. On the other hand, pushing an element should result in the size of the stack to increase by one.

A QuickCheck test case execution starts from the initial state and selects a transition whose precondition is enabled. The action upon the transition is actuated and the postcondition is checked. If a

[1] `http://www.cs.waikato.ac.nz/~marku/mbt/modeljunit/`

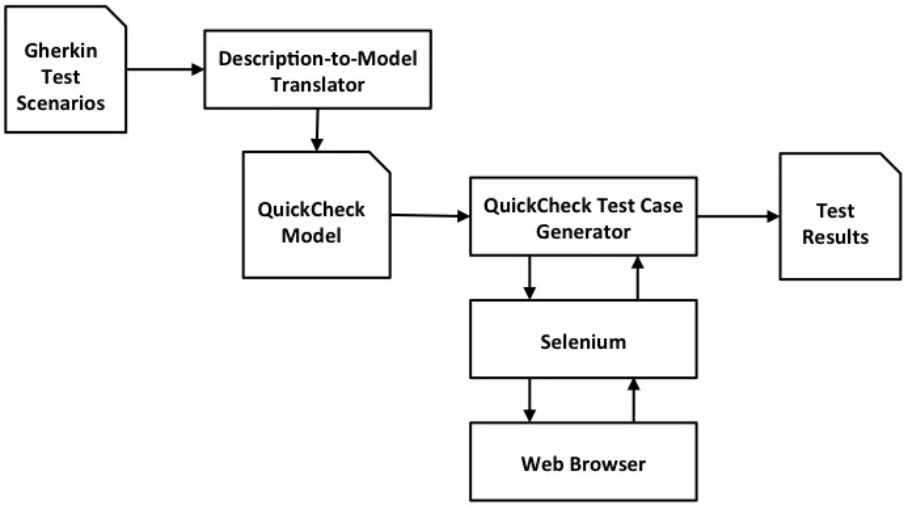

Figure 4: An overview of the system architecture

postcondition check fails, then this signifies a failed test.

Once QuickCheck has been chosen to instantiate our approach, a tool was developed as illustrated in the architecture diagram in Figure 4. The system takes a set of test scenarios as input and proceeds to parse them and convert them into a QuickCheck model. This model is then fed to QuickCheck for test case generation and execution. During the test case execution process, Selenium is utilised for interacting with a web browser and exercising the web application. Finally, QuickCheck outputs a report containing results of the test run.

Having implemented an instantiation of the technique, we could explore its applicability through an industry case study. This is the subject of the following section.

5 Case Study

Our initial evaluation of the technique presented here is a case study. The objectives of this case study were threefold:

1. Demonstrate the feasibility of the technique on a non-trivial system
2. Compare the technique to its manual alternative
3. Note any relevant observations to help form future research directions

The case study was carried out on myHealth[2], a web-based solution that has recently been developed for the government of Malta's Ministry of Health to provide an improved and more efficient service to patients in the Maltese health care infrastructure. It consists of a web portal allowing access to myHealth data for registered citizens. A back-end portal is provided so as to allow monitoring and management to administrators. The case study was scoped to the front-end portal only, mainly because the back-end portal was purchased off the shelf and no test scenarios were available for it. The system is able to distinguish between doctors and citizens, providing different functionality for the two roles with Table 1 providing a list of high level functionality available to both types of users. When citizens gain access to

[2]`http://www.myhealth.gov.mt`

Analyse existing test scripts
Write tests in Gherkin format
Create automation code for steps
Generate Quickcheck model
Run model-based tests

Figure 5: The process followed when carrying out the case study.

their myHealth Record account, they are able to search for doctors listed in the Medical Council Register, request a specific doctor to be their *trusted doctor* and go on to view and manage all their personal health data.

The case study was carried out as depicted in Figure 5. Since the system's test scripts were not documented in Gherkin format, they needed to be analysed and written in the required form. Once this was done, test automation code was created such that all the scenarios became executable automated tests. This enabled us to feed the test scenarios into our tool, create QuickCheck models and subsequently generate and execute tests based on the models. These steps are elaborated in the following subsections.

5.1 Constructing a model for the case study

When we were granted access, the myHealth Record system had already been through testing and been launched to the public. In this section we briefly outline the testing process used by the company that developed myHealth. The system went through two phases of testing: the first phase was a unit testing and integration phase which was mainly carried out by developers in the company. The second phase consisted of user acceptance testing and was carried out in conjunction with the customer in order to verify that the required functions had been implemented and were working correctly. The nature of the work presented in this paper is mainly concerned with system-level or user acceptance testing, so the second phase is of interest to us. This phase was driven by pre-documented test cases which outlined a series of procedures and expected outcomes (see Table 2 for a sample test case).

For the purposes of this case study, the test cases from the second phase of testing were secured and used as the basis to construct specifications in the Gherkin language. Each test case was manually analysed and systematically converted to its equivalent *Given-When-Then* format. Furthermore, step definitions were implemented such that each scenario became an executable automated test. Related test scenarios were then grouped together and models were generated accordingly using an implementation of Algorithm 1. Figure 6 shows one of the models generated during the exercise.

5.2 Results and Observations

With regards to exploring the feasibility of such an approach on a non-trivial system, the case study demonstrated that this is indeed the case. We were able to generate models directly from business level specifications and subsequently utilise those models to generate and execute test cases. The process took approximately seven days of work to complete. This was similar to the length of time taken by the company to carry out manual user acceptance testing of the product. However, if one looks closer at how the time was invested (see Table 3), in the case of manual testing, the company estimates that four days were spent designing test cases whilst three days were spent executing them. In the case of automated model-based testing, the test design aspect was practically free (due to the use of QuickCheck) whilst the actual execution only took one day. Whilst on the face of it, it seems that both approaches are equally time consuming, it is worth noting that if the company had adopted the use of Cucumber

Actions available to doctors	**Actions available to citizens**
• Login as a doctor • Go to *appointments* page – Search for a patient's appointments – View a patient's appointments • Go to *case summaries* page – Search for a patient's case summary – View a patient's case summary – Print patient's case summary • Go to *laboratory results* page – Search for a patient's lab result – View a patient's lab result – Release result to patient – Mark lab results as read or unread – Print patient's lab results • Go to *medical image reports* page – Search for a patient's report – View a patient's report – Release report to patient – Mark medical image reports as read or unread – Print patient's medical image report • Go to *POYC entitlement* page – Search for a patient – View a patient's entitlement • Go to *search patients' data* page – Search for a patient – View a patient's records • Go to *manage patients* page – Search for a patient – Remove a patient • Go to *manage notifications* page – Subscribe to email notifications	• Go to *appointments* page – View appointments • Go to *case summaries* page – Search for case summary – View a case summary – Print case summary • Go to *laboratory results* page – Search for a lab result – View a lab result – Mark lab results as read or unread – Print lab results • Go to *medical image reports* page – Search for a report – View a report – Mark medical image report as read or unread – Print medical image reports • Go to *POYC entitlement* page – View a POYC entitlement – Print entitelment • Go to *search for doctors* page – View all registered doctors – Search for a doctor – Send request to a doctor • Go to *manage doctors* page – Remove a doctor • Go to *manage notifications* page – Subscribe to email notifications – Subscribe to mobile notifications

Table 1: Functionality available to the doctors and citizens in the myHealth system.

Test Case Number:	4
Scope:	Releasing the full Updated Report through the Patient Data menu
Procedure and Expected Outcome:	
Doctor - Releasing the report 1. Login to the system as Doctor - e-ID 0240860M 2. In the Doctor Landing Page, open the *Patient Data* accordion 3. Navigate to the *Medical Image Reports* section in the accordion 4. Locate the previous report and Click on *View Reports* 5. Click *Release All* 6. Logout Citizen - Testing that the updated report has been released 1. Login to the system as Citizen - e-ID 0240851M 2. In the Citizen Landing Page, open the *myMedical Image Reports accordion*. Note that the *Updated box* in the accordion shows 1 3. Locate a report and note that an *Updated Reports* icon is displayed 4. Click on *View Reports* and note that some of the reports were updated. Previous reports should be marked as *Superseded* and greyed out while the updated report should be marked as *Final (Updated)*. The report should contain the updates. 5. Logout	
Date Run:	01/03/2012
Run By:	<Names of testers present>
Overall Test Outcome:	Pass
Additional Commands and/or Screenshots:	

Table 2: A sample test case from the user acceptance testing phase

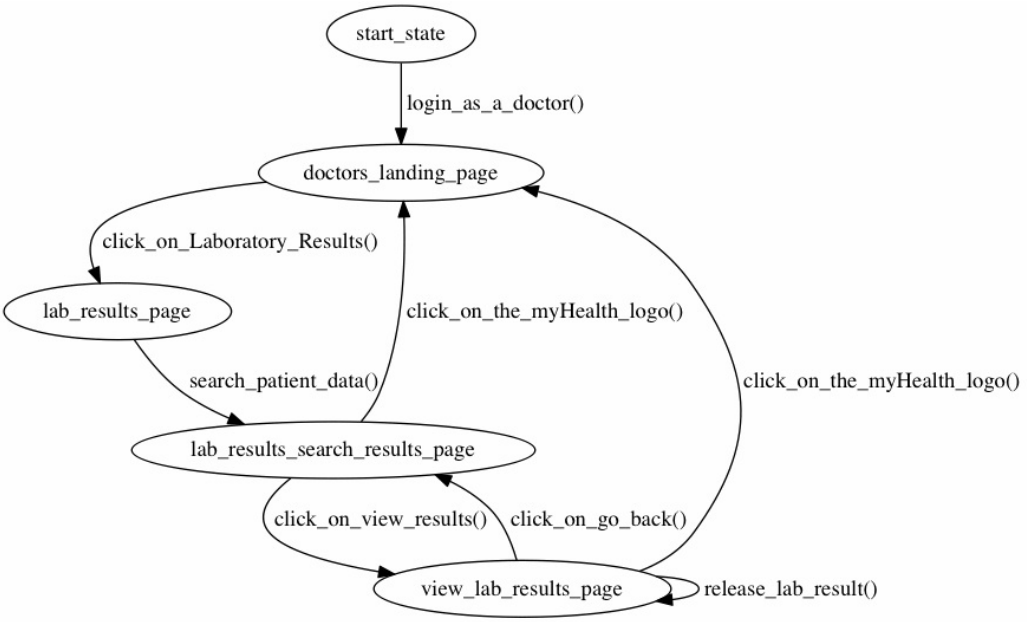

Figure 6: Model generated for myHealth's *Doctor's Lab Results* functionality

	Manual Testing	**Automated Model-based Testing**
Generation of test cases	4 days	-
Creating and automation Gherkin Scenarios	-	6 days
Execution of test cases	3 days	1 day
Total time	7 days	7 days

Table 3: Comparison of the time required for manual testing and automated model-based testing.

in its normal operations (as is the case with an increasing number of companies), the time require for automated model-based testing would have been substantially reduced. Furthermore, automated model-based testing is efficiently repeatable and will deliver consistent results when compared to that delivered manually by human testers.

Two main observations were made throughout this process: Firstly, automated testing will not be able to replicate all tests which are carried out manually by a human tester. In fact, we were able to replicate 74% of manual tests due to test cases which either required human judgement (e.g. checking printouts) or were too complex to feasibly automate (e.g. temporal tests). The second observation was related to the shrinking feature of QuickCheck: while this particular feature is fast in console-style applications, we observed that shrinking took an inordinately long amount of time when testing web applications. This is mainly due to the expense required to repeatedly launch a browser and navigate through a series of web pages during the shrinking process.

It is worth noting that no new bugs were uncovered as a result of applying the technique. We think that this is due to the fact that by the time we ran the case study, the system had been thoroughly tested, deployed and in active use for quite some time.

5.3 Threats to Validity

The evaluation presented here is mostly concerned with demonstrating that the automated generation of models from natural language specifications and the subsequent generation and execution of tests is prima facie feasible. Consequently, the case study is limited in terms of depth and is subject to a number of external validity threats. We split these threats into two categories: (i) threats related to new challenges introduced by this technique and (ii) threats related to test engineering challenges which are manifested in most automation attempts. We argue that the first category of threats are highly plausible and should be investigated further whilst the second category of threats is diminished due to the fact that we are reusing existing automation code which should have already dealt with them.

The main threat introduced by the current approach is concerned with the requirement that conventions be followed when constructing natural language specifications. That is to say that preconditions, actions and post conditions are assumed to reside in specific elements within each scenario. Whilst the conventions are arguably reasonable, they do introduce some restrictions on what is meant to be a flexible language. This constitutes a threat to validity because unlike this case study whereby scenarios were created for the purpose of the evaluation, real-life scenarios will be written by different people who may or may not have followed the same conventions.

With regards to test engineering challenges, two representative issues come to mind. Firstly, one can also argue that the data used in this case study was not sufficiently complex and in reality, data dependencies may be highly complex. Whilst the concern is justified, we argue that the threat to validity here is limited, mostly due to the fact that the technique assumes that an automated testing infrastructure already exists. Such infrastructures usually contain sophisticated test data factories which support automated testing. These factories are typically called by '*Given*' steps in order to create/fetch any required data as well as by '*Then*' steps to tear down and residual test data that is no longer required. Similar arguments can be made for other classes of test engineering challenges. For example, some test could be sensitive to timing issues. However, since we are essentially reusing existing test automation code which is exposed to the same challenges, it can be safely assumed that test engineers would have solved these challenges within that code. If not, then the original set of automated test would not be functional in the first place. Whilst this should reduce validity concerns related to data dependencies, we do intend to observe this in more elaborate case studies as discussed in Section 7.

6 Related Work

The application of model based testing to web applications is not new and has been proposed on numerous occasions. Andrews et al. [1] proposed the use of finite state machines as an ideal model for this domain. Marchetto et al. [11] extended this approach to cater for more dynamic Ajax-based applications. Similarly, Wu and Offutt [15] propose a technique to model both static and dynamic aspects of web applications based on identifying elements of dynamically created web pages and combining them using sequence, selection, aggregation and regular expressions to create composite web pages. Ernits et al. [7] demonstrate how model based testing can be applied to web applications using the NUnit tool. NUnit allows developers to create models in C# and subsequently uses those models to generate test cases. Whilst the authors had success with the technique, they commented that a significant amount of time was spent learning the modelling formalism and building the test harness [7]. This contrasts to our approach which bypasses the need of manually programming the models as long as the business-level specifications adhere to the conventions highlighted in this paper.

With regards to automatic model generation, the literature mostly reveals work utilising reverse engineering approaches. Hajiabadi and Kahani [9] propose a technique for first modelling the structural aspects of a web application by observing client requests and server responses, and subsequently applying ontologies to generate test data. The authors claim that ontologies are required in order to make test case selection more effective. Our approach differs in that ours focus on the generation of sequences of actions rather than the data used. (This is left as future work in our case.) Ricca and Tonella [12] propose a UML representation of web applications and present a tool called *ReWeb* which reverse engineers web applications into this notation. They also present a tool named *TestWeb* which processes UML representations of web sites to automatically generate and execute test cases. While this approach is similar to ours in the use of a high-level specification notation, it differs in that it uses the implementation of the system itself to extract the model while we use the high-level business specification to this end.

The closest work to ours is that carried out by Jääskeläinen et al. [10] which similarly synthesises and merges test models from simple linear test cases. The main difference is that our propositions are not state-based but rather transition-based, resulting in simpler states. The implication is that our merging is fully automatable (with the possibility of having duplicate transitions) while in their case the merging has to be manually checked.

7 Conclusions and Future Work

In this paper we proposed an approach for converting business level specifications written using *Given-When-Then* natural language notation into models that can subsequently be used to generate and execute test cases. Whilst the literature contains multiple proposals for model-based testing of web applications, most techniques require the involvement of highly technical individuals who were capable of constructing models of their systems. Given the dynamic evolutionary nature of modern systems, maintaining models in the long run may not be desirable or even feasible. Our technique can work off natural language specifications which adhere to three simple conventions as discussed in Section 3 and if applied in a company where automated testing is already the norm, model creation and subsequent test case generation can be achieved at minimal cost.

With regards to future work, there are a number of paths which we would like to pursue:

- The first involves converting the technology to make use of different model-based testing tools on various platforms. This is desirable because the most expensive part of the approach is the automation of test scenarios in the language of the model-based testing tool, which in our case was Erlang. If models can be generated in the same language as that being used by developers in a particular company, then test automation code can simply be reused by models out of the box.
- The second path of research involves carrying out larger case studies both in terms of system size, as well as in terms of project duration. It would be ideal to find industry partners who have ongoing projects using Cucumber as a test automation tool and would like to utilise our technique on a long term basis. This would help us make more detailed observations and modify the technique accordingly.
- While our approach is a perfect fit for web applications due to the ease with which scenarios can be connected over common web pages, we are confident the approach can also be applied to other domains such as desktop GUI-based applications where the connecting element might be a tab or a window which is enabled at a particular point in time.

- On a more theoretic level, we wish to explore the possibility of modelling variations of data elements rather than limiting ourselves to sequences of actions. For example in the case study we could model different laboratory results, doctor attributes, patient conditions, etc., automatically to generate more varied test cases.
- Finally, it is rarely the case that scenarios run in a sequential fashion in a real life environment. Thus, it would also be interesting to explore models which allow modelling several users concurrently executing different scenarios.

References

[1] Anneliese A Andrews, Jeff Offutt & Roger T Alexander (2005): *Testing web applications by modeling with FSMs*. *Software & Systems Modeling* 4(3), pp. 326–345, doi:10.1007/s10270-004-0077-7.

[2] Thomas Arts, John Hughes, Joakim Johansson & Ulf Wiger (2006): *Testing telecoms software with quviq quickcheck*. In: *Proceedings of the 2006 ACM SIGPLAN workshop on Erlang*, ACM, pp. 2–10, doi:10.1145/1159789.1159792.

[3] David Chelimsky, Dave Astels, Bryan Helmkamp, Dan North, Zach Dennis & Aslak Hellesoy (2010): *The RSpec Book: Behaviour Driven Development with Rspec, Cucumber, and Friends (The Facets of Ruby Series)*. Pragmatic Bookshelf.

[4] Koen Claessen & John Hughes (2011): *QuickCheck: a lightweight tool for random testing of Haskell programs*. *Acm sigplan notices* 46(4), pp. 53–64, doi:10.1145/357766.351266.

[5] Alistair Cockburn & Jim Highsmith (2001): *Agile software development, the people factor*. *Computer* 34(11), pp. 131–133, doi:10.1109/2.963450.

[6] Ian Dees, Matt Wynne & Aslak Hellesoy (2013): *Cucumber Recipes: Automate Anything with BDD Tools and Techniques*. Pragmatic Bookshelf.

[7] Juhan Ernits, Rivo Roo, Jonathan Jacky & Margus Veanes (2009): *Model-based testing of web applications using NModel*. In: *Testing of Software and Communication Systems*, Springer, pp. 211–216, doi:10.1007/978-3-642-05031-2_14.

[8] Gherkin: *Gherkin Wiki*. Available at `http://github.com/cucumber/cucumber/wiki/Gherkin`.

[9] Hamideh Hajiabadi & Mohsen Kahani (2011): *An automated model based approach to test web application using ontology*. In: *Open Systems (ICOS), 2011 IEEE Conference on*, IEEE, pp. 348–353, doi:10.1109/ICOS.2011.6079282.

[10] Antti Jääskeläinen, Antti Kervinen, Mika Katara, Antti Valmari & Heikki Virtanen (2009): *Synthesizing Test Models from Test Cases*. In: *Hardware and Software: Verification and Testing*, *LNCS* 5394, Springer, pp. 179–193, doi:10.1007/978-3-642-01702-5_18.

[11] Alessandro Marchetto, Paolo Tonella & Filippo Ricca (2008): *State-based testing of Ajax web applications*. In: *Software Testing, Verification, and Validation, 2008 1st International Conference on*, IEEE, pp. 121–130, doi:10.1109/ICST.2008.22.

[12] Filippo Ricca & Paolo Tonella (2001): *Analysis and testing of web applications*. In: *Proceedings of the 23rd international conference on Software engineering*, IEEE Computer Society, pp. 25–34, doi:10.1109/ICSE.2001.919078.

[13] Selenium: *Selenium Project Page*. Available at `http://code.google.com/p/selenium`.

[14] Specflow: *Specflow - Pragmatic BDD for .NET*. Available at `http://www.specflow.org`.

[15] Ye Wu & Jeff Offutt (2002): *Modeling and testing web-based applications*. Technical Report, George Mason University. GMU ISE Technical ISE-TR-02-08.